UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/751,090	12/29/2000	Richard W. Busser	4430-22	5019

22442 7590 09/02/2004

SHERIDAN ROSS PC
1560 BROADWAY
SUITE 1200
DENVER, CO 80202

EXAMINER	
MASKULINSKI, MICHAEL C	

ART UNIT	PAPER NUMBER
2113	

DATE MAILED: 09/02/2004

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	09/751,090	BUSSER, RICHARD W.	
	Examiner	Art Unit	
	Michael C Maskulinski	2113	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *15 July 2004*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-20* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☒ Claim(s) *20* is/are allowed.

6)☒ Claim(s) *1-14 and 16-19* is/are rejected.

7)☒ Claim(s) *15* is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other: _____.

Non-Final Office Action

Claim Rejections - 35 USC § 112

1. In view of the recent amendments the rejection of claims 1, 3-8, 10, 11, and 13 under 35 U.S.C. 112, first paragraph have been withdrawn.

Claim Rejections - 35 USC § 103

2. The text of those sections of Title 35, U.S. Code not included in this action can be found in a prior Office action.

3. Claims 1-8, 10-12, 14, 16-19 are rejected under 35 U.S.C. 103(a) as being unpatentable over Renner, Jr., U.S. Patent 6,243,827 B1, and further in view of Jeffries et al., U.S. Patent 5,974,544.

Referring to claim 1:

a. In column 6, lines 26-36, Renner, Jr. discloses a RAID 3 system comprising a single host computer, a RAID controller, and two tiers of 5 Direct Access Storage Device (DASD) units with two parity DASDs and two additional hot spare DASDs (using firstly an array of storage devices to conduct read/write operations under control of at least a first controller including providing metadata to each of said storage devices of said array).

b. In column 12, lines 13-15, Renner, Jr. discloses a MEDIUM ERROR status reported by the storage unit. Further in column 1, lines 53-59, Renner, Jr. discloses that medium errors are transient errors such as bus errors (ascertaining that a failure has occurred).

c. In column 4, lines 30-47, Renner, Jr. discloses that catastrophic disk-array failures involve the failure of greater than one disk in a RAID 1-5 system, or any disk in a RAID 0 system (discontinuing use of at least two storage devices of said array related to conducting read/write operations based on the failure).

d. In column 11, lines 48-54, Renner, Jr. discloses that if the data was not successfully committed to the array even after retries and/or removal of a redundant disk from operation, the subroutine ADD BAD BLOCKS is called for the range of logical addresses described by the command (using said array of storage devices after said discontinuing).

e. Renner, Jr. doesn't explicitly disclose that said using said array of storage devices after said discontinuing includes updating said metadata to remove an indication that none of said at least two storage devices are accessible and to indicate that said at least two storage devices are accessible, wherein said updating said metadata is performed in response to said ascertaining that a failure has occurred, and wherein said updating said metadata is performed without interrupting power to the array. In column 16, lines 49-56, Jeffries et al. disclose that on power up, DDA determines which drives have failed, which drives have failed previously and now seem OK, and whether a drive has been replaced. DDA can differentiate between new drives and previously bad drives, by looking at the DDA sector; new drives are assumed not to have a DDA sector. To reuse a drive, which has previously failed, the diagnostics must be used to erase the drive's DDA sector. It would have been obvious to one of ordinary skill

at the time of the invention to include Jeffries et al. ability to reuse a drive into the system of Renner, Jr. A person of ordinary skill in the art would have been motivated to make the modification because having the ability to recover a previously failed drive saves the cost and time of having to replace the drive with a new one. It is not cost effective to replace a drive that is not failed.

Referring to claim 2, in column 2, lines 3-11, Renner, Jr. discloses transient failures and controller failure. In column 4, lines 53-56, Renner Jr. discloses a power failure (a transient failure related to a back plane). In column 1, lines 53-59, Renner, Jr. discloses that medium errors are transient errors such as bus errors (a cable).

Referring to claim 3, in column 4, lines 30-34, Renner, Jr. discloses that catastrophic disk-array failures involve the failure of greater than one disk in a RAID 1-5 system, or any disk in a RAID 0 system. In almost all cases, however, there remains the ability of the controller to communicate with at least one disk in the array (said using said array of storage devices after said discontinuing includes making a determination related to being able to use said array of storage devices including said at least two thereof).

Referring to claim 4, in column 11, lines 31-36, Renner, Jr. discloses the process, which is performed when a write operation to the storage array has been completed, successfully or unsuccessfully, by the controller. In the case of a failure, this process will only be called after the number of retries prescribed by the error handling policies of the controller have been performed (said making said determination related to be able

to use said array of storage devices including said at least two thereof includes checking whether one or more of said storage devices is off-line).

Referring to claim 5, in column 4, lines 34-39, Renner, Jr. discloses that the present invention uses software and a small portion of each disk in the array to write a bad area table on each disk. The bad area table provides the logical address and length of the area in the array's logical space, which has been corrupted by physical damage on the media or other causes of write failure (said updating said metadata includes modifying metadata in a primary dead partition map on each of said storage devices, wherein each of said storage devices including said at least two storage devices is indicated as being valid).

Referring to claims 6 and 18, in column 9, lines 31-59, Renner, Jr. discloses that if BBM STAMP is valid, its timestamp is checked against the timestamp of CANONICAL STAMP (where the timestamp for a null stamp is defined to be older than any valid timestamp). If it is more recent than CANONICAL STAMP, CANONICAL STAMP is set to BBM STAMP. Control then passes to the next iteration of the loop. When there are no more disks to check, the CANONICAL STAMP is checked for a null value. If it is not null, then an appropriate stamp has been located, and all bad regions described in CANONICAL STAMP are added to the BBM MAP. Further, in column 16, lines 54-56, Jeffries et al. disclose that to reuse a drive which has previously failed, the diagnostics must be used to erase the drive's DDA sector (said using said array of storage devices after discontinuing includes issuing a trust array command to said first controller that

causes said updating said metadata including writing all zeros in a primary dead partition map).

Referring to claim 7, in column 9, lines 31-59, Renner, Jr. discloses that the presence of other disks is checked in a loop (said using said array of storage devices after said discontinuing includes determining whether each of said storage devices of said array is accessible after said issuing of said trust array command).

Referring to claim 8, in column 11, lines 36-40, Renner, Jr. discloses that if the array has redundancy and errors occur on only one disk, the disk may be removed from operation prior to the invocation of this process, in which case the status of the write operation will be considered to be good (said using said array of storage devices after said discontinuing includes controlling re-use of said array when it is determined that no more than one of said storage devices of said array is off-line).

Referring to claim 10, in column 11, lines 36-40, Renner, Jr. discloses that if the array has redundancy and errors occur on only one disk, the disk may be removed from operation prior to the invocation of this process, in which case the status of the write operation will be considered to be good (said using said array of storage devices after said discontinuing includes controlling re-use of said array based on one of a user determination and an automatic determination independently of the user).

Referring to claim 11, in column 6, lines 59-67 continued in column 7, lines 1-6, Renner, Jr. discloses that the flow of data between host and disk array is indicated by the heavy line. Data is received from the host computer via the host SCSI bus into the SCSI input/output processor (SCSI IOP). The SCSI IOP initiates memory transactions

to or from the cache memory through the bridge chip, which bridges the system bus and the cache bus. A cache bus connects the bridge chip, cache memory, and the hardware control mechanism DMA Sync. The DMA Sync acts as a direct memory access (DMA) controller with the additional RAID-3 function of parity generation and checking and replacement of data with a hot spare. It also generates reads or writes to specific cache addresses and translates the data between the cache bus and the SCSI interface chip on the individual channel (said using said array of storage devices after said discontinuing includes generating a command by a host and transmitting said command to said first controller).

Referring to claim 12, it is inherent to have the read/write command, as discussed above, initiated manually by the user of the array.

Referring to claim 14:

a. In column 6, lines 26-36, Renner, Jr. discloses a RAID 3 system comprising a single host computer, a RAID controller, and two tiers of 5 Direct Access Storage Device (DASD) units with two parity DASDs and two additional hot spare DASDs (an array of storage devices relative to which read and write data transfers are conducted).

b. In column 6, lines 26-36, Renner, Jr. discloses a RAID 3 system comprising a single host computer, a RAID controller, and two tiers of 5 Direct Access Storage Device (DASD) units with two parity DASDs and two additional hot spare DASDs (a controller communicating with said array of storage devices for conducting read/write operations).

c. In column 6, lines 59-67 continued in column 7, lines 1-6, Renner, Jr. discloses that the flow of data between host and disk array is indicated by the heavy line. Data is received from the host computer via the host SCSI bus into the SCSI input/output processor (SCSI IOP). The SCSI IOP initiates memory transactions to or from the cache memory through the bridge chip, which bridges the system bus and the cache bus. A cache bus connects the bridge chip, cache memory, and the hardware control mechanism DMA Sync. The DMA Sync acts as a direct memory access (DMA) controller with the additional RAID-3 function of parity generation and checking and replacement of data with a hot spare. It also generates reads or writes to specific cache addresses and translates the data between the cache bus and the SCSI interface chip on the individual channel (a host communicating with said controller that makes requests related to data to be stored and data to be obtained from said array of storage devices).

d. In column 4, lines 30-47, Renner, Jr. discloses that catastrophic disk-array failures involve the failure of greater than one disk in a RAID 1-5 system, or any disk in a RAID 0 system. In almost all cases, however, there remains the ability of the controller to communicate with at least one disk in the array. The present invention uses software and a small portion of each disk in the array to write a bad area table on each disk. The bad area table provides the logical address and length of the area in the array's logical space, which has been corrupted by physical damage on the media or other causes of write failure. After a catastrophic failure of multiple disks, assuming at least one disk can be written

to, there will be a record of the failure on at least one disk (wherein said host is used in generating a trust array command related to updating metadata on each of said storage devices of said array after a fault occurs and after use of said array was discontinued due to the fault).

f. Renner, Jr. doesn't explicitly disclose changing the metadata from indicating that said one or more storage devices is inaccessible to indicating that said one or more storage devices is accessible without restarting said array, and wherein said trust array command is generated in response to detection of said fault. In column 16, lines 49-56, Jeffries et al. disclose that on power up, DDA determines which drives have failed, which drives have failed previously and now seem OK, and whether a drive has been replaced. DDA can differentiate between new drives and previously bad drives, by looking at the DDA sector; new drives are assumed not to have a DDA sector. To reuse a drive, which has previously failed, the diagnostics must be used to erase the drive's DDA sector. It would have been obvious to one of ordinary skill at the time of the invention to include Jeffries et al. ability to reuse a drive and indicate its status into the system of Renner, Jr. A person of ordinary skill in the art would have been motivated to make the modification because having the ability to recover a previously failed drive saves the cost and time of having to replace the drive with a new one. It is not cost effective to replace a drive that is not failed.

Referring to claim 16, in column 4, lines 30-47, Renner, Jr. discloses that catastrophic disk-array failures involve the failure of greater than one disk in a RAID 1-5

system, or any disk in a RAID 0 system. In almost all cases, however, there remains the ability of the controller to communicate with at least one disk in the array. The present invention uses software and a small portion of each disk in the array to write a bad area table on each disk. The bad area table provides the logical address and length of the area in the array's logical space, which has been corrupted by physical damage on the media or other causes of write failure. After a catastrophic failure of multiple disks, assuming at least one disk can be written to, there will be a record of the failure on at least one disk. The task of repairing the array is greatly simplified because all bad regions of the array can be easily identified (said trust array command is generated independently of any reconstruction and/or restoration of said array).

Referring to claim 17, in column 11, lines 31-40, Renner, Jr. discloses that in the case of a failure, if the array has redundancy and errors occur on only one disk, the disk may be removed from operation prior to the invocation of this process (said host controls discontinuing use of said array of storage devices based on the fault). Further, in column 11, lines 48-54, Renner, Jr. discloses that if the data was not successfully committed to the array even after retries and/or the removal of a redundant disk from operation, the subroutine ADD BAD BLOCKS is called for the range of logical addresses described by the command (and subsequently said host receives an input that is used in generating said trust array command).

Referring to claim 19, in column 11, lines 31-36, Renner, Jr. discloses that in the case of a failure, the process will only be called after the number of retries prescribed by the error handling policies of the controller have been performed (a determination is

made by at least one of said controller and said host related to whether one or more of said storage devices is off-line before said trust array command is generated). Further, in column 4, lines 30-47, Renner, Jr. discloses that after a catastrophic failure of multiple disks and the bad area table is written, the system determines which disks are accessible (a determination is made by at least one of said host and said controller related to whether each of said storage devices of said array is accessible after said trust array command is generated).

4. Claim 13 is rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of Renner, Jr., U.S. Patent 6,243,827 B1 and Jeffries et al., U.S. Patent 5,974,544 as applied to claim 1 above, and further in view of Jones, U.S. Patent 5,479,653.

Referring to claim 13, neither Renner, Jr. nor Jeffries et al. explicitly disclose not restoring and/or reconstructing user data and/or parity. In column 3, lines 16-20, Jones discloses that as failures occur, the disk array system automatically reconfigures one or more of the drives to other RAID configurations which utilize less data redundancy and thus require a lesser number of drives. It would have been obvious to one of ordinary skill at the time of the invention to include the disk array reconfiguration of Jones into the combined system of Renner, Jr. and Jeffries et al. A person of ordinary skill in the art would have been motivated to make the modification because *the disk array system initially takes advantage of all of the available drives for the maximum performance and redundancy, thus maintaining a high level of fault tolerance* (see Jones: column 3, lines 22-25).

5. Claim 9 is rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of Renner, Jr., U.S. Patent 6,243,827 B1 and Jeffries et al., U.S. Patent 5,974,544 as applied to claim 1 above, and further in view of Stephenson, U.S. Patent 6,353,895 B1.

Referring to claim 9, in the Abstract, Renner, Jr. discloses a software-based method for facilitating the recovery of a RAID storage system from the simultaneous failure of two or more disks. However, neither Renner, Jr. nor Jeffries et al. explicitly disclose allowing data and/or parity to be read by said first controller when more than one of said storage devices is off-line and reading said data and/or parity from said storage devices of said array that are on-line. In column 2, lines 37-49, Stephenson discloses that one embodiment includes a redundant array of independent disk drives that provides one-drive and two-drive fault tolerance. Data recovery from a one or two drive failure is accomplished by using a two-dimensional XOR parity arrangement. It would have been obvious to one of ordinary skill at the time of the invention to include the two-drive fault tolerance and parity arrangement of Stephenson into the combined system of Renner, Jr and Jeffries et al. A person of ordinary skill in the art would have been motivated to make the modification because *this parity arrangement uses less storage than mirroring when the number of total drives is greater than four* (see Stephenson: Abstract).

Allowable Subject Matter

6. Claim 15 is objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

7. Claim 20 is allowed.

Response to Arguments

8. Applicant's arguments filed July 15, 2004 have been fully considered but they are not persuasive.

9. On page 9, under the section REMARKS, the Applicant argues, "On power up (emphasis by Applicant), the method discussed by Jeffries determines which drives have failed, which drives have failed previously and now seem OK, and whether a drive has been replaced. (Jeffries, col. 16, lns. 49-51). Accordingly, the rewriting of metadata in response to detecting a failure, without interrupting power to the array, is not described by Jeffries." Further, on page 10, under the section REMARKS, the Applicant argues, "In particular, the cited references do not teach, suggest or disclose reusing storage devices after a failure, or rewriting metadata to allow access to such storage devices in response to ascertaining that a failure has occurred and without interrupting power to the array as recited by Claim 1. The Examiner respectfully disagrees. In column 10, lines 54-56, Jeffries et al. disclose **To reuse a drive which has previously failed** (emphasis by Examiner), the diagnostics must be used to erase the drive's DDA sector. Further, in column 16, lines 40-41, Jeffries et al. disclose that information (metadata) is stored on all the other disks, and so it is maintained across power up cycles and resets. In other words, when a reset or power up occurs the information is

used to indicate which drives have failed, which drives have failed previously and now seem OK, and whether a drive has been replaced. Nowhere do Jeffries et al. state that a restart or power up is necessary to update the metadata.

10. On page 11, under the section REMARKS, the Applicant argues, "the Jeffries reference does not teach, suggest or disclose changing metadata without restarting the array, or generating a trust array command in response to detecting a fault." The Examiner respectfully disagrees for at least the reasons in paragraph 9 above.

11. Applicant's arguments on pages 11-12, under the section REMARKS, with respect to claims 16-19 fail to comply with 37 CFR 1.111(b) because they amount to a general allegation that the claims define a patentable invention without specifically pointing out how the language of the claims patentably distinguishes them from the references.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Michael C Maskulinski whose telephone number is (703) 308-6674. The examiner can normally be reached on Monday-Friday 9:30-6:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Robert W Beausoliel can be reached on (703) 305-9713. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

MM

ROBERT BEAUSOLIEL
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100